


IFW

PTO/SB/21 (04-04)

Approved for use through 07/31/2006. OMB 0651-0031

Patent and Trademark Office: U.S. DEPARTMENT OF COMMERCE

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.



TRANSMITTAL FORM

(to be used for all correspondence after initial filing)

Application Number 10/711,320

Filing Date 9/10/04

First Named Inventor Rene Glaise

Art Unit 2661

Examiner Name

Total Number of Pages in This Submission

Attorney Docket Number FR920030066US1

ENCLOSURES (check all that apply)

<input type="checkbox"/> Fee Transmittal Form	<input type="checkbox"/> Drawing(s)	<input type="checkbox"/> After Allowance communication to Technology Center (TC)
<input type="checkbox"/> Fee Attached	<input type="checkbox"/> Licensing-related Papers	<input type="checkbox"/> Appeal Communication to Board of Appeals and Interferences
<input type="checkbox"/> Amendment / Reply	<input type="checkbox"/> Petition	<input type="checkbox"/> Appeal Communication to TC (Appeal Notice, Brief, Reply Brief)
<input type="checkbox"/> After Final	<input type="checkbox"/> Petition to Convert a Provisional Application	<input type="checkbox"/> Proprietary Information
<input type="checkbox"/> Affidavits/declaration(s)	<input type="checkbox"/> Power of Attorney, Revocation Change of Correspondence Address	<input type="checkbox"/> Status Letter
<input type="checkbox"/> Extension of Time Request	<input type="checkbox"/> Terminal Disclaimer	<input type="checkbox"/> Other Enclosure(s) (please identify below):
<input type="checkbox"/> Express Abandonment Request	<input type="checkbox"/> Request for Refund	
<input type="checkbox"/> Information Disclosure Statement	<input type="checkbox"/> CD, Number of CD(s) _____	
<input checked="" type="checkbox"/> Certified Copy of Priority Document(s)	Remarks	
<input type="checkbox"/> Response to Missing Parts/ Incomplete Application		
<input type="checkbox"/> Response to Missing Parts under 37 CFR 1.52 or 1.53		

SIGNATURE OF APPLICANT, ATTORNEY, OR AGENT

Firm or Individual name	Joseph P. Abate, Reg. No. 30,238
Signature	
Date	December 15, 2004

CERTIFICATE OF TRANSMISSION/MAILING


I hereby certify that this correspondence is being facsimile transmitted to the USPTO or deposited with the United States Postal Service with sufficient postage as first class mail in an envelope addressed to: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450 on the

Typed or printed name	Teri McDonald
Signature	
Date	December 15, 2004

This collection of information is required by 37 CFR 1.5. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to 12 minutes to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.

If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.

THIS PAGE BLANK (USPTO)



**Europäisches
Patentamt**

**European
Patent Office**

**Office européen
des brevets**

Bescheinigung

Certificate

Attestation

Die angehefteten Unterlagen stimmen mit der ursprünglich eingereichten Fassung der auf dem nächsten Blatt bezeichneten europäischen Patentanmeldung überein.

The attached documents are exact copies of the European patent application described on the following page, as originally filed.

Les documents fixés à cette attestation sont conformes à la version initialement déposée de la demande de brevet européen spécifiée à la page suivante.

Patentanmeldung Nr. Patent application No. Demande de brevet n°

03368090.1

Der Präsident des Europäischen Patentamts:
Im Auftrag

For the President of the European Patent Office

Le Président de l'Office européen des brevets
p.o.

R C van Dijk

THIS PAGE BLANK (USPTO)



Anmeldung Nr:
Application no.: 03368090.1
Demande no:

Anmeldetag:
Date of filing: 18.09.03
Date de dépôt:

Anmelder/Applicant(s)/Demandeur(s):

INTERNATIONAL BUSINESS MACHINES CORPORATION

Armonk, NY 10504
ETATS-UNIS D'AMERIQUE

Bezeichnung der Erfindung/Title of the invention/Titre de l'invention:
(Falls die Bezeichnung der Erfindung nicht angegeben ist, siehe Beschreibung.
If no title is shown please refer to the description.
Si aucun titre n'est indiqué se référer à la description.)

Method and system to enable an adaptive load balancing in a parallel packet
switch

In Anspruch genommene Priorität(en) / Priority(ies) claimed /Priorité(s)
revendiquée(s)

Staat/Tag/Aktenzeichen/State/Date/File no./Pays/Date/Numéro de dépôt:

Internationale Patentklassifikation/International Patent Classification/
Classification internationale des brevets:

H04L12/56

Am Anmeldetag benannte Vertragstaaten/Contracting states designated at date of
filing/Etats contractants désignées lors du dépôt:

AT BE BG CH CY CZ DE DK EE ES FI FR GB GR HU IE IT LU MC NL
PT RO SE SI SK TR LI

THIS PAGE BLANK (USPTO)

**METHOD AND SYSTEM TO ENABLE AN ADAPTIVE
LOAD BALANCING
IN A PARALLEL PACKET SWITCH**

Field of the Invention

5 The present invention relates to high speed switching of
data packets in general and, is more particularly concerned
with a method and a system to adapt the load balancing of the
incoming traffic over the planes of a parallel packet switch
(PPS) on the basis of the monitoring of requests and acknowl-
10 edgments exchanged between ingress port adapters and arrays of
collapsed virtual output queues (cVOQ) situated within the
plane switch cores.

Background of the Invention

To design and implement a high capacity packet-switch typically, having throughputs ranging from a few hundreds giga (10⁹) to multi tera (10¹²) bits per second and aggregating traffic from/to OC-768c or 40 Gbps ports, a practical architecture often considered to overcome the technology limitations that prevent from reaching directly this level of performance, is a parallel packet switch (PPS) architecture. As sketched in figure 1, a PPS is thus comprised of multiple identical lower-speed packet-switch planes e.g., four planes (100), operating independently and in parallel. In each ingress port-adapter (IA), such as (110), an incoming flow of packets (120) is thus spread (130), packet-by-packet, by a load balancer across the slower packet-switches, then recombined by a multiplexing function (140) in the egress part of each port-adapter (EA) e.g., (150). As seen by an arriving packet, a PPS is a single-stage packet-switch that needs to have only a fraction of the performance necessary to sustain a PPS port data rate (125). If four planes (100) are for example used, as shown in figure 1, their input (102) and output (104) ports need only to have one fourth of the performance that would otherwise be required to handle a full port data rate. More specifically, as an example, the four independent switches shown can have OC-192c or 10Gbps ports. When associated they deliver OC-768c or 40 Gbps port speed provided that ingress and egress port-adapters (110, 150) are able to transparently load balance and recombine the packets. This approach is well known from the art and sometimes referred to as 'Inverse Multiplexing' or 'load balancing'. Among many publications on the subject one may e.g., refer to a paper published in Proc. ICC'92, 311.1.1-311.1.5, 1992, by T. ARAMAKI et al., untitled 'Parallel "ATOM" Switch Architecture for High-Speed ATM Networks' which discusses the kind of architecture considered here.

The above scheme is also very attractive because of its inherent capability to support redundancy. By placing more

planes than what is strictly necessary thus, five in the example of figure 1, it is possible to hot replace a defective plane without having to stop traffic. When a plane is detected as being, or becoming, defective ingress adapter load balancers
5 can skip the defective plane and e.g., switch to a spare plane waiting in hot standby. When all the traffic from the defective plane has been drained out it can be removed and replaced by a new one. Hence, a PPS structure potentially also offers a graceful mode of degradation since, if installed planes are all
10 used, one can still disable one of them and keep traffic moving through the ones left even though this is at a lower level of performance.

Thus, if PPS is really attractive to permit the implementation of very high-performance switches, with an architecture
15 which is intrinsically redundant, a key requirement of carrier-class switches which need to operate 24-hour a day and 7-day a week, it remains that this is only achievable if load balancing (130) in each IA can actually performed accordingly. Especially, IA's must indeed be able to detect that a switching plane does
20 not operate properly so as it can all be skipped. And, ideally, load balancing should be able to detect that a single port is not functioning properly in order that only traffic destined for it be not load balanced through the corresponding (partially) defective plane thus, fully taking advantage of a
25 graceful mode of degradation since only one port on one plane would need to be skipped in that case while the rest of the switch continues to operate adequately.

Finally, load balancing should better be self-adaptive too so as it can dynamically react to the actual behavior of the
30 switching planes. Although all IA's load balancing functions will attempt to evenly spread the traffic over all available planes and ports, hot spots may occasionally occur for a while in certain planes due to the nature of the traffic. Adaptive load balancing can take care of this by dynamically loading
35 more the planes that are not congested.

It is then the purpose of the present invention to disclose a method and a system which enable a self adaptation of IA load balancing for all the situations mentioned above namely, to skip a complete switching plane if found defective
5 or just one or a few ports of it, if a plane is only partially defective.

Also, self adaptation must occur if a plane, or ports of a plane, are detected as being temporarily congested in which case load balancing must tend to avoid loading more the
10 congested plane and ports while congestion is in effect.

Object of the Invention

Thus, it is a broad object of the invention to remedy the shortcomings of the prior art as described here above.

It is another object of the invention to provide a system
15 and a method for dynamically evaluating the relative level of occupancy of the individual switching planes of a parallel packet switch.

It is a further object of the invention to provide a system and a method for dynamically adapting load balancing by
20 analyzing in real time the actual behaving of the switching planes, without having to increase the exchanged of flow control information.

The accomplishment of these and other related objects is achieved by a method in an ingress port-adapter (IA) for
25 dynamically evaluating the relative level of occupancy of the individual switching planes of a parallel packet switch, wherein a request is transmitted from said IA to one of said individual switching planes each time a data packet is received in said IA, said data packet being temporarily stored in said
30 IA, and wherein an acknowledgment is returned from said one of

said individual switching planes to said IA when said data packet can be processed, said method comprising the steps of:

computing, for each individual plane, the number of data packets waiting to be processed;

5 determining the range of data packets waiting to be processed among all individual switching planes; and,

comparing said range of data packets waiting to be processed with at least one predetermined threshold,

10 signaling each at least one predetermined threshold exceeded.

Brief Description of the Drawings

Figure 1 represents a well known parallel packet switch architecture, comprising multiple identical lower-speed packet-switch planes operating independently and in parallel, wherein incoming flow of packets are spread in each ingress port-adaptor, packet-by-packet, by a load balancer across the slower packet-switches, and then recombined in the egress port-adapters.

Figure 2 describes the kind of parallel packet switch architecture adapted to implement the invention.

Figure 3 discusses how plane congestion can be evaluated from each ingress port-adaptor that can count independently the number of requests sent to each plane of a parallel packet switch structure.

Figure 4 further elaborates on the behaving of the used up/down counters.

Figure 5 describes how counters values must be used to

determine when a resource is malfunctioning thus, should be skipped or replaced, and how load balancing over resources can be adapted.

Detailed Description of the Preferred Embodiment

The invention assumes that PPS switching planes are of the kind shown in Figure 2, referred to as collapsed virtual output queuing array (cVOQ). The mode of operation of such a fabric assumes that a request is sent from IA's (200) to the switch core (210) each time a data packet is received in an IA's, the data packet being temporarily stored in the IA. The switch core (210) returns acknowledgments (240) if there are enough switching resources left in switch core to allow incoming packets (205), waiting in IA's queues (215), to be admitted in core. For each incoming packet (205) there is a request (207) sent to the switch core so as to update, within switch core, the image of the corresponding IA queue, here taking the form of a simple counter (262) of requests. In this particular example of a switch core allowing to carry out the invention, switching is achieved with a shared-memory (212). Packets waiting in input queues (220) are admitted on reception of an acknowledgment (240) issued if enough room is left in shared-memory and if there are not too many packets already waiting for the corresponding output port, e.g., (255). The selection of which acknowledgments are returned is done on the basis of an algorithm that examines the image of all IA's queues (260) and referred to as a collapsed virtual output queuing or cVOQ array. Hence, cVOQ is in switch core the image of all IA's queues organized per destination (225), and in the general case per class of service, implementing a scheme well-known from the art referred to as 'virtual output queuing' (VOQ) which prevents all form of 'head of line' (HoL) blocking in IA's.

Whichever type of switch core is actually used the invention rests on the assumption that there is a means of evaluating the occupancy of the switch core, as viewed from each IA, for the particular traffic it received. Here, this is simply
5 done by monitoring the requests versus the returned acknowledgments. Hence, if there is a persistent or growing deficit of acknowledgments as compared to the forwarded requests this is indicative that the plane is getting congested. This is further discussed in next figure.

10 **Figure 3** discusses how plane congestion can be evaluated from each IA (310) that can count independently the number of requests sent (330) to each plane (300) of a PPS structure. Counting is achieved with a set of up/down counters (315) incremented with each request departing from IA. It is decre-
15 mented each time an acknowledgment (305) is returned from the corresponding plane.

Eventually, counters (315) must return to zero when all packets have been forwarded which is indicative of the fact that they all have been indeed processed. However, on a steady
20 state traffic, the up/down counters normally indicate a deficit of acknowledgments since there are always some requests waiting in switch core to be acknowledged or traveling back to the source (320). Similarly, if switch planes are some distance apart thus connected through cables to the port adapters, as
25 shown in figure 2 (250), the requests do not reach immediately the switch core. Depending on the physical implementation of a PPS structure and taking into consideration the processing time of the electronic boards that carry out the switching functions, round trip (325) time (RTT) may be significant and
30 has to be expressed in packet-times at port speeds considered by the invention i.e., 10 Gbps and above.

If above implicitly assumes there is one up/down counter (315) per plane in order to count the overall number of requests load balanced per plane, nothing prevents from having

a counter dedicated to each output port (340) though. This permits to obtain a finer monitoring of the hardware resources down to a per link and per plane basis. Hence, if the counter associated to a port of a particular plane starts to shift up
5 while the others (the equivalent ones on the other planes) stay balanced the problem can be pinpointed to a particular link that ingress load balancer must start to skip enabling an automatic graceful mode of degradation. This should be reported to a central control point in order to schedule a maintenance
10 or raise some form of alert.

Similarly, if all the counters for a particular destination (output port) start to shift together this clearly means that the whole egress adapter i.e., (150) in figure 1, is no longer functioning properly, or the communications line(s)
15 connected to it. This may trigger another type of alert and maintenance.

Table of figure 3 (350) summarizes the above discussion. Obviously, contingent to a particular implementation, more refinements could be brought to the isolation and graceful
20 degradation modes possible thus, affecting only part of the hardware resources while the rest is still up and running. Also, the reporting to a central control point of all the information collected independently in IA's can help to further pinpoint a type of failure.

25 Figure 4 further elaborates on the behaving of the up/down counters. Again, there must have one up/down counter per physical resource to be monitored. This can be e.g., a whole switching plane, each port of a switching function, each port of a plane, down to each high-speed serial link of which port planes
30 are made. Whichever level of granularity is considered, there must have requests (410) and acknowledgments (420) that are uniquely associate-able to a particular up/down counter (400) thus, identifying a given resource that will be possibly skipped if found defective, so as to support a graceful mode of

degradation, or replaced by an equivalent resource held in standby.

Therefore, each counter (400) reflects the current difference between the forwarded requests (402) minus the returned
5 acknowledgments (404). For a full continuous utilization of the resource associated to the counter this latter normally reaches a steady-state count of RTT packets (406) since one acknowledgment is coming back, after RTT packet-times, for each forwarded requests. Hence, at establishment of the traffic, value of the
10 counter normally grows (430) and stabilizes (435) to RTT when resource is fully utilized. Obviously, any intermediate value (440) up to RTT, is likely to be observed over time since, resource may not be actually fully utilized and traffic is generally expected to fluctuate between zero and full utilization of the resource until it may returned eventually to zero
15 (450) when no longer utilized.

However, if resource is malfunctioning, or is prevented from operating at nominal speed, counter value is going to shift up eventually beyond RTT (445) because requests must
20 still be issued for processing new arriving packets, received from an upward situated source of traffic, while there is no, or less than nominal, return of acknowledgments. Therefore, the up shifting of a counter resource may be as well the result of an action taken by a downward device to flow-control its
25 upwards situated counterpart. Networks, where switch fabrics are utilized, indeed need to implement a flow-control mechanism to avoid congestion and to pace the flows of traffic. This takes generally the form of a backpressure type of signal (475) issued by a downward device (470) to its upward source(s) of
30 traffic here a switch according to the invention (460). If congestion lasts, this latter must in turn, issue back-pressure signals (465) to the actual upward source(s) of traffic such as (480). Hence, the chief reason for which an up/down counter monitoring a resource may shift up beyond RTT is the result of
35 actions taken by the network flow-control mechanism when a

downward situated device cannot accept more traffic or at a pace lower than nominal. Then, counter (400) has associated thresholds (406) that can be used to determine when back-pressure signal (465) must be propagated.

5 Therefore, it must be clear that a counter such as (400) cannot be taken alone to decide if the particular resource it monitors is malfunctioning or not. Depending on the actual level of utilization of a network and of a switch fabric such as (460), the prime reason for which up/down counters may shift
10 up beyond their normal RTT value is more likely due to network flow-control attempting to pace the flows of traffic especially, when traffic reaches, or is close to, full utilization of a set of available resources. Hence, groups of counters must rather be compared to decide if an individual resource,
15 part of a group, is indeed malfunctioning.

At this point it is important to notice that RTT must be expressed in packet-time units of the corresponding resource to which a particular up/down counter is associated. In the type of switch fabric considered by the invention traffic is essentially load balanced on many parallel hardware resources so as
20 to be able to reach the expected level of performance and/or to provide for redundancy. As an example, a switch fabric having ports specified at 40 Gbps actually, having to operate at 80 Gbps in order to deliver a speedup factor of two to take care
25 of protocol overheads (e.g., encapsulation and segmentation and reassembly) and to absorb bursts of traffic may require 4 planes with 20-Gbps ports. In turn, implementing a 20-Gbps port will require the actual use of 8 high-speed serial links at 2.5 Gbps effective data rate. Hence, a typical 64-byte (i.e.:
30 512-bit) packet time when expressed in Ns (10^{-9} second) is only $512/80 = 6.4$ Ns for a 80-Gbps port. However, packets are actually load balanced over four planes in this example. For an individual port plane, packet duration is four times higher i.e.: $512/20 = 25.6$ Ns . And for an individual link at 2.5 Gbps,
35 over which packets are finally load balanced so as to allow

redundancy and graceful degradation down to a link, packet duration becomes $512/2.5 = 205$ Ns. Hence, for a counter associated to a link, the rate at which packet are forwarded and received is, at full speed, of one per 205 Ns in this particular example. RTT must then be expressed in this unit of time for a counter watching a single link.

Figure 5 describes how counters values must be used to determine when a resource is malfunctioning thus, should be skipped or replaced, and how load balancing over resources can be adapted.

Up/down counters of the invention are assumed to belong to a group (510). Such a group is e.g., the group (310) of four counters shown in figure 3. In this particular elementary example there is one counter (515) per switching plane such as (300) which are the resources to be overall monitored. The switching planes are assumed to be identical thus, are expected to behave identically so that individual counter values can be compared (545, 535) together to make decisions. Because planes are identical, IA's load balancing function i.e., (130) in figure 1, is also assumed to load identically the four planes. Therefore, malfunction of one of the resources, here a switching plane, is determined through a different behaving of one of the counters vs. the others. If there is a single counter that grossly shifts up (518) while the others stay balanced (545) this is the clear indication that the corresponding plane is no longer functioning properly. This is true even if network flow-control discussed in figure 4 has issued back-pressure signals. In this case, all the planes are normally affected identically thus, all counters shift together. Indeed, if an egress switch port adapter such as (150) in figure 1 is getting congested because it is no longer authorize to forward traffic to a downward device e.g., (470) of figure 4, the resulting backpressure is eventually applied to all switching planes (100) that, together, implement a switch port.

Therefore, this is the comparison of the counters (520) rather than their individual values that permit to decide when a resource does no longer operate as expected.

5 An alert (530) can thus be issued to a control device when the result of the counters comparison (520) shows that one of them (518) is grossly or significantly behaving differently. Depending on a particular application, values beyond which an alert must be triggered can have to be set differently. Also, alert can be issued only after a significant period of time in
10 order not to trigger any false alert. Setting of the thresholds, and integration times to trigger an alert, can therefore be very different from one application of the invention to another. Especially, threshold setting can be carried out in many ways. In its simplest form it can be an absolute
15 range of waiting packets i.e., the observed difference between the minimum and maximum number of packets waiting in individual planes (535), which must not be normally exceeded. It can be as well a dynamic threshold expressed e.g., as a relative proportion of the total number of waiting packets in the switching
20 planes or any other criterion applicable to a particular implementation of the invention that allows to pinpoint on the misbehaving of a particular switching resource.

Finally, when counters are all overall balanced (545) the smaller differences between their values can be further
25 analyzed so as to feed back the information to the ingress adapters which can better adapt load balancing (540). As an example, statistical behaving of individual counters can be evaluated and their average values and/or standard deviations measured and compared from which decisions can be derived to
30 better adapt load balancing over the switching planes. Clearly, if it can be determined that a counter has an average value which is always, or statistically significantly below the others then load balancing should attempt to load more the corresponding plane in proportion of the observed difference.

Above are just examples of what can be achieved with the up/down counters according to the invention. Those skilled in the art will recognize that many other applications can be considered beyond their utilization to monitor the resources of
5 a switch fabric so as to implement all forms of active and passive redundancy or to adapt dynamically load balancing in a PPS like structure.

THIS PAGE BLANK (USPTO)

Claims:

What is claimed is:

1. A method in an ingress port-adapter (IA) for dynamically evaluating the relative level of occupancy of the individual switching planes of a parallel packet switch, wherein a request is transmitted from said IA to one of said individual switching planes each time a data packet is received in said IA, said data packet being temporarily stored in said IA, and wherein an acknowledgment is returned from said one of said individual switching planes to said IA when said data packet can be processed, said method comprising the steps of:

computing, for each individual plane, the number of data packets waiting to be processed;

determining the range of data packets waiting to be processed among all individual switching planes; and,

comparing said range of data packets waiting to be processed with at least one predetermined threshold,

signaling each at least one predetermined threshold exceeded.

2. The method according to claim 1 wherein said computing step is automatically performed by an up/down counter incremented with each said transmitted request and decremented with each said returned acknowledgment.

3. The method according to either claim 1 or claim 2 wherein said signaling step further comprises the step of:

determining which at least one switching plane is responsible for said at least one predetermined threshold exceeded.

4. The method according to claim 3 wherein said signaling step and said determining steps are used for triggering the further step of:

adapting load balancing over said switching planes.

5 5. The method according to claim 4 wherein said adapting step consists in reducing load balancing over said at least one determined switching plane.

6. The method according to claim 4 wherein said adapting step consists in skipping said at least one determined switching
10 plane by said load balancing.

7. The method according to claims 1 to 6 wherein said signaling step execution is contingent to a minimum duration of a threshold crossing.

8. The method according to claims 1 to 7 wherein said signaling
15 step is used to raise an alert signal depending on which said at least one threshold is exceeded.

9. The method according to claims 1 to 8 wherein said predetermined threshold is an absolute number of packets waiting to be processed.

20 10. The method according to claims 1 to 8 wherein said predetermined threshold is a relative number of packets waiting to be processed.

11. The method according to any one of the previous claims wherein the level of occupancy of the switching resources are
25 evaluated beyond said switching planes down to any identifiable physical switching resource and load balancing performed accordingly.

12.The method according to claim 11 wherein said identifiable physical switching resource is an individual serial link.

13.An apparatus comprising means adapted for carrying out each step of the method according to any one of the claims 1 to 12.

5 14.A computer-like readable medium comprising instructions for carrying out each step of the method according to any one of the claims 1 to 12.

THIS PAGE BLANK (USPTO)

**METHOD AND SYSTEM TO ENABLE AN ADAPTIVE
LOAD BALANCING
IN A PARALLEL PACKET SWITCH**

Abstract

5 A method and a system to adapt the load balancing of the incoming traffic over the planes of a parallel packet switch (PPS) on the basis of the monitoring of requests and acknowledgments exchanged between ingress port adapters and arrays of collapsed virtual output queues (cVOQ) situated within the
10 plane switch cores is disclosed. According to the invention, at least one counter is associated, in each ingress port-adaptor, to each individual switching plane or device to be monitored. Each of these counters is incremented when a request is sent to the corresponding individual switching
15 plane or device and decremented when an acknowledgment is received from this individual switching plane or device. When the range of values taken by the counters of a same ingress port-adaptor reaches a predetermined threshold, less (or none) incoming traffic is further transmitted to the individual
20 switching plane or device associated to the higher value counter. An alarm signal is possibly raised too e.g., for replacing the defective individual switching plane or device.

Figure 3.

THIS PAGE BLANK (USPTO)




Figure 1





Figure 2



350 Up/Down Counters Shifting Up:	
One for a particular port of a plane	One link is out of order
All for a particular port	One egress adapter, or the communications line(s) connected to it, is out of order
All for a particular plane	One switching plane is out of order

Figure 3




Figure 4

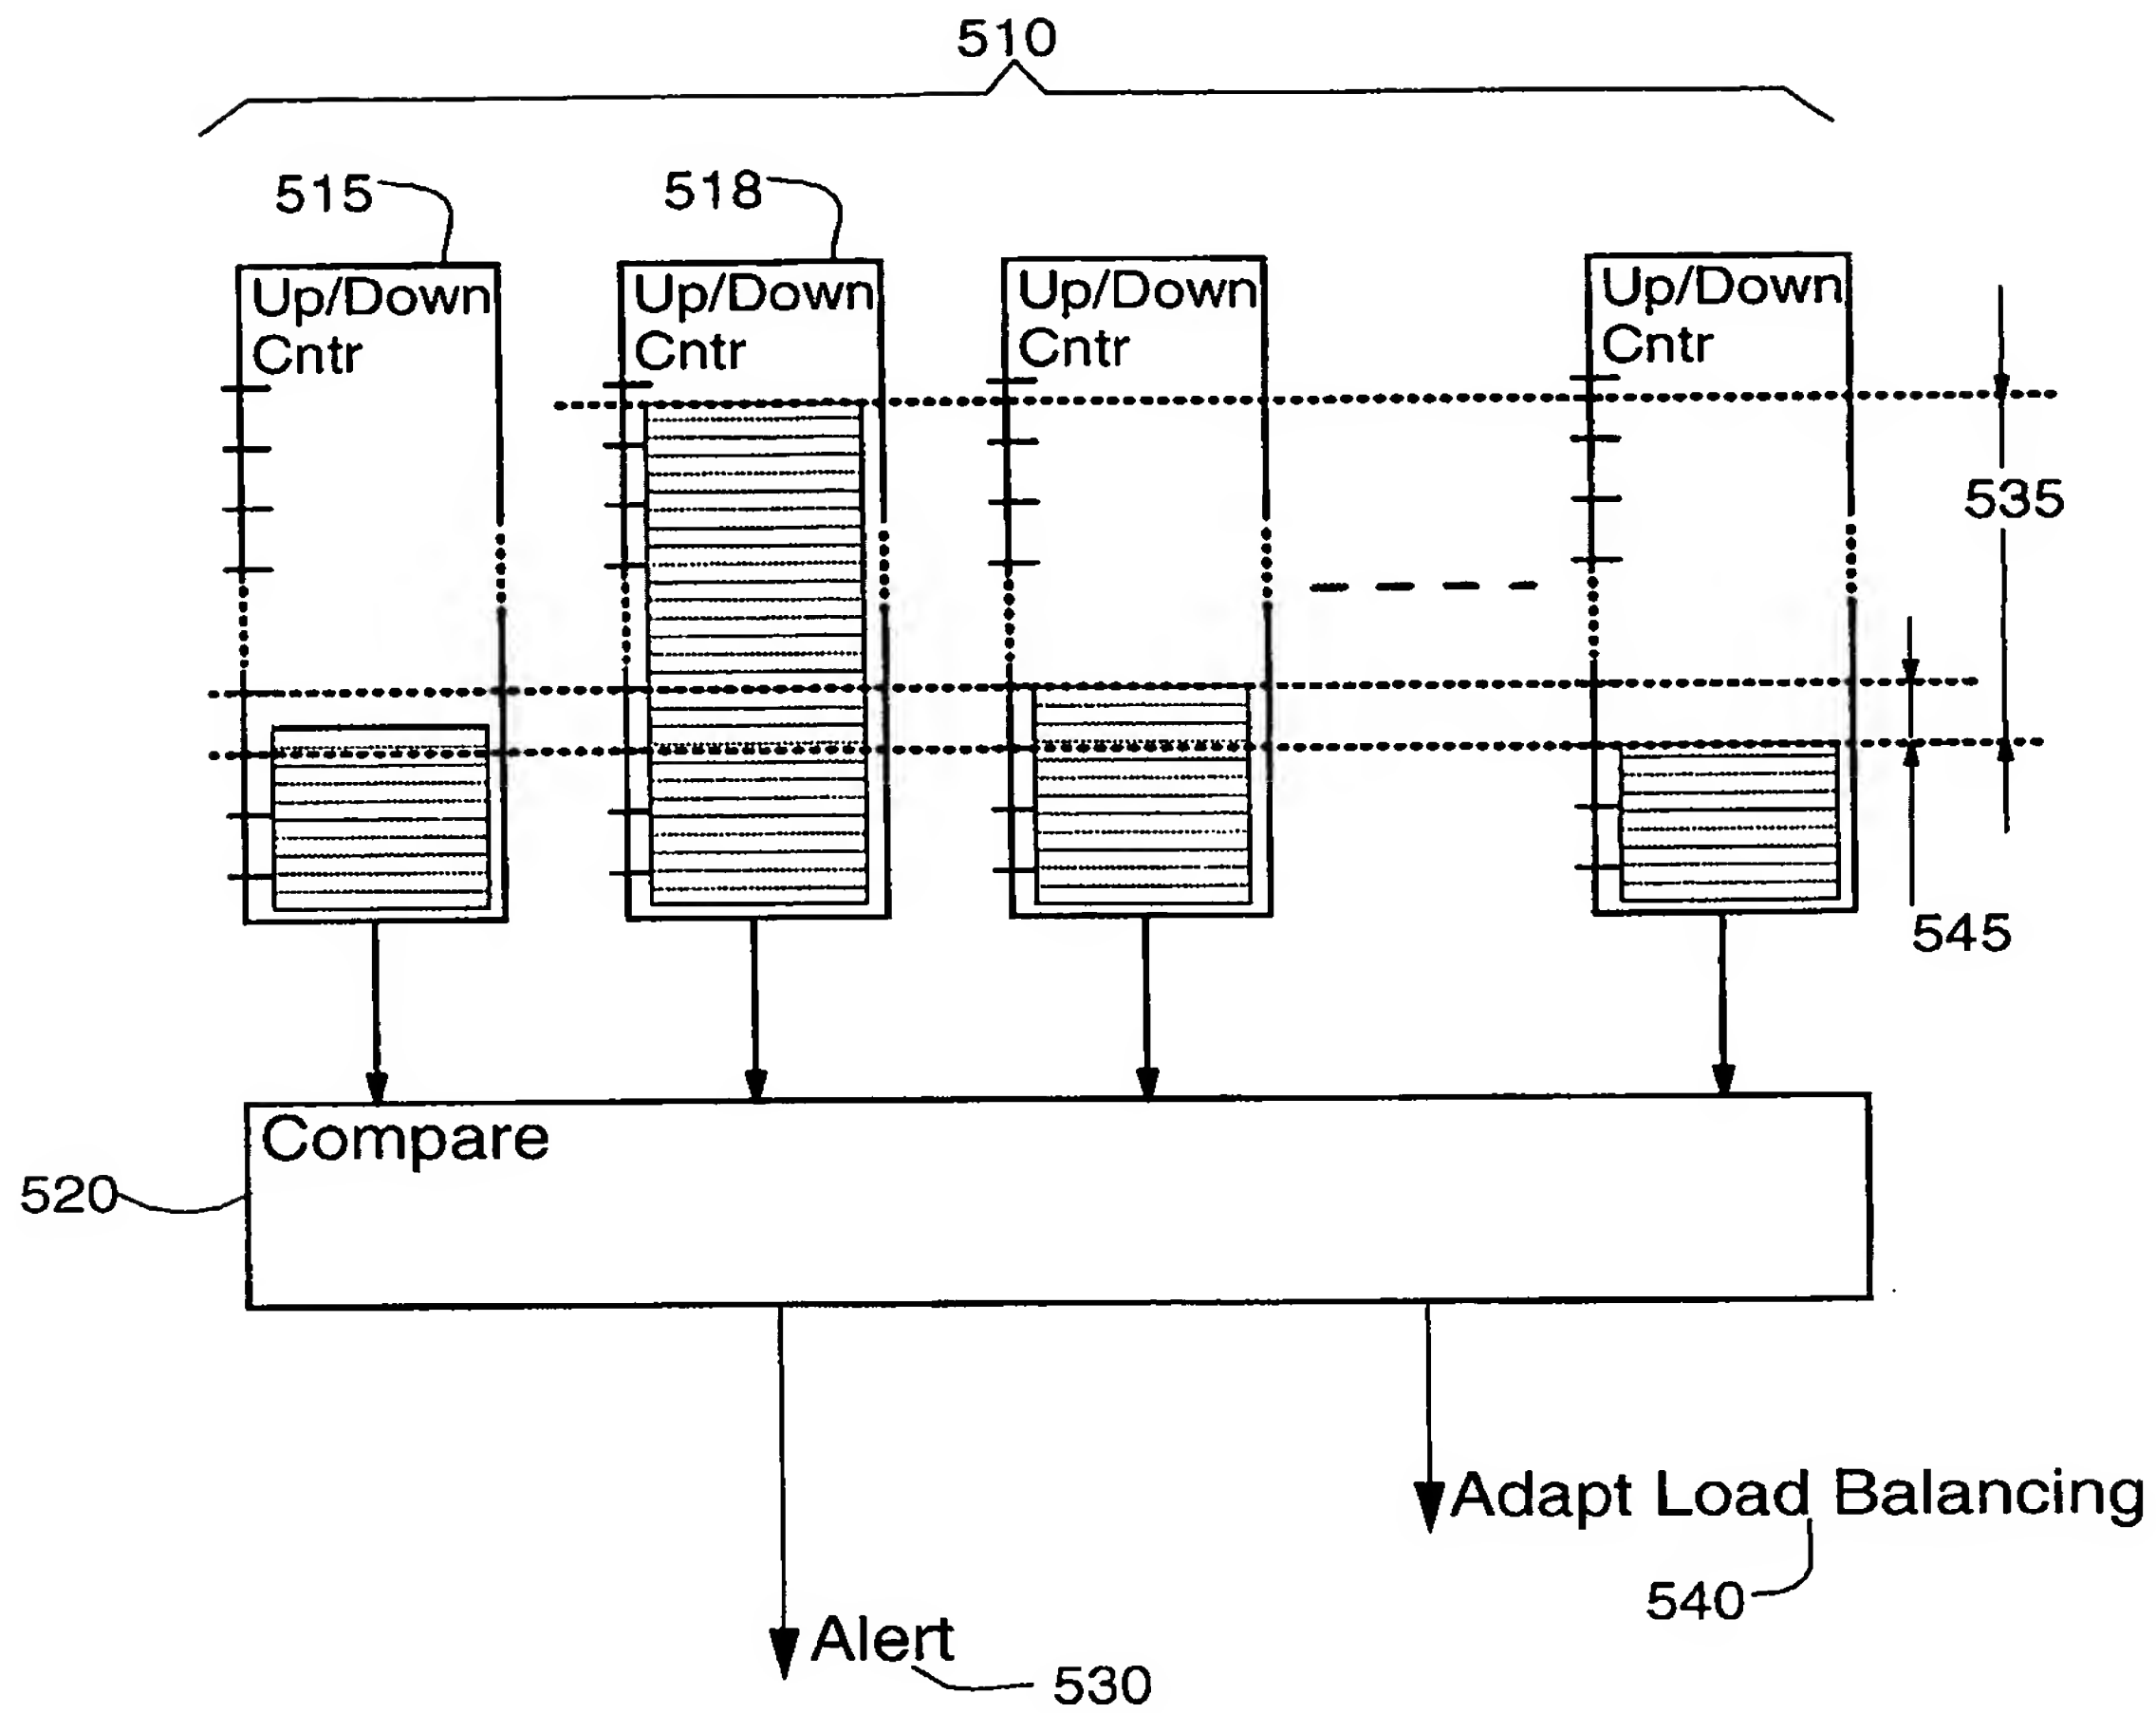


Figure 5

THIS PAGE BLANK (USPTO)